



By Jason Anderson
Product Line Manager,
Eaton

Multi-tenant data center trend: distributed power in the white space

Executive summary

Cloud computing vendors and multi-tenant colocation facilities make every effort to maximize the scalability, efficiency and agility of their data centers. Unlike private data centers, their power configurations are closely scrutinized by their customers and prospects who are comparing to them other service options in a very competitive environment. As a result, more and more of them are investigating replacing older, transformer-based centralized power protection schemes with distributed architectures in which uninterruptible power systems (UPSs) reside in the white space. This white paper describes the advantages of positioning UPSs in the white space as well as the essential qualities to look for in a white space-ready UPS.

Table of contents

Advantages of putting UPSs in your white space	2
Characteristics to look for in a white space-ready UPS	3
1. Low TCO	3
2. Flexible deployment	3
3. Superior manageability	4
Conclusion	4
About Eaton	4
About the author	4



Powering Business Worldwide

Advantages of putting UPSs in your white space

Multi-tenant data center (MTDC) operators and cloud operators have long appreciated the merits (scalability and ease of installation) of distributed power protection schemes, in which UPSs reside in close proximity to server, storage and networking equipment. Until recently, however, most incumbent UPSs which were scaled to the needs of enterprise data centers were too hot, large and inflexible for use in the white space, where cooling and floor space are precious commodities and mechanical, electrical and plumbing components can impose significant deployment constraints.

Now, however, more efficient, compact and adaptable technologies have revolutionized both the centralized and distributed UPS applications. Modularity, inherent redundancy, and ultra-high efficiency now characterize the best large UPS systems. These same attributes have migrated into smaller, even more compact systems, making distributed power protection a practical option for companies large and small. Clients (users) whose rapid growth rate has exceeded the infrastructure capacity in the gray space, are ideal candidates for a white space UPS solution. Other applications include those datacenters which are deploying high-performance, high-density computing on a segmented basis. These users need highly redundant, 2N type, scalable systems, to be deployed in a small area of the white space, but may plan to expand in repeatable, identical sections (like colo cages), requiring a compact dedicated UPS. Cloud computing vendors and colocation data center operators are increasingly taking advantage of that development, and collecting the following benefits as a result:

- **Rapid flexibility and scalability** – Many MTDC operators offer multiple levels of redundancy, with some customers choosing 2N and others N+1, for example. Additionally, cloud service providers grow at unpredictable rates that can spike sharply with little warning. As a result, they usually optimize their data centers for swift expansion. Distributed power schemes are a good fit for such environments as they feature smaller, lighter UPSs that take less time and effort to install than the older, heavier and bulkier models typically used in centralized power designs. Moreover, distributed power architectures scale far less disruptively than centralized schemes, as organizations in need of more capacity can simply install additional smaller UPSs instead of taking their entire data center offline while they replace a sole large UPS with a new and even bigger one.
- **Higher flexibility and scalability** – Many MTDC operators offer multiple levels of redundancy, with some customers choosing 2N and others N+1, for example. Additionally, cloud service providers grow at unpredictable rates that can spike sharply with little warning. As a result, they usually optimize their data centers for swift expansion. Distributed power schemes are a good fit for such environments as they feature smaller, lighter UPSs that take less time and effort to install than the older, heavier and bulkier models typically used in centralized power designs. Moreover, distributed power architectures scale far less disruptively than centralized schemes, as organizations in need of more capacity can simply install additional smaller UPSs instead of taking their entire data center offline while they replace a sole large UPS with a new and even bigger one.

- **Higher availability** – MTDC operators are now used to potential customers asking about how they architect to remove single points of failure. Reducing the distance between UPSs and the loads they protect shortens and simplifies the power chain, resulting in fewer potential points of failure. For example, distributed architectures often don't require centralized downstream static switches. That means fewer devices for administrators to maintain and fewer opportunities for malfunctioning hardware to create downtime. What's more, in a distributed power design each UPS protects only a small portion of a company's servers, so if a problem occurs between the UPS and its downstream loads little capacity is lost. In centralized architectures, power chain failures can impact every server in the white space.
- **Lower capital expenses** – Old style, oversized power architectures may lead companies to make large upfront investments in older technology UPSs. Modern centralized and distributed power schemes, by contrast, enable organizations to buy only as much capacity as they need to meet early demand, and then cost-effectively add further capacity incrementally as their needs increase over time. This is particularly important in a multi-tenant environment where customers have new requirements for expansion within the same environment and prospective customers also bring very specific requirements to a selection process.
- **Greater energy efficiency** – Cloud and multi-tenant providers alike struggle to accurately forecast operations needs, as those requirements are dependent upon the growth of their business. Deploying more UPS capacity than your data center initially requires wastes power as well as capital. Distributed power designs allow you to raise energy efficiency by matching capacity to demand with greater precision, providing a critical business benefit for a provider.

A note about distributing 480V power

Organizations with distributed power schemes can reduce capital expenses even further by utilizing UPSs with 480V input/output voltage instead of the 208/120V power common in most U.S. data centers today. The cables used to connect 480V UPSs to a data center's utility input service are thinner, lighter and therefore less expensive than the cables they must use with 208V UPSs. In addition, unlike 208V UPSs, 480V units don't require a neutral wire, producing further savings by enabling companies to use three wires instead of four.

Of course, most IT equipment in the U.S. requires 208V/120V power, so companies with distributed architectures featuring 480V UPSs must step power down between their UPSs and servers. However, a conventional power distribution unit or a matching cabinet installed next to the UPS cabinet can perform that function without significantly increasing upfront expenditures.

Characteristics to look for in a white space-ready UPS

Capitalizing on the benefits of distributed power protection schemes takes UPS hardware with an extensive and rigorous set of capabilities that fall within three categories: low total cost of ownership (TCO), high deployment flexibility and strong manageability.

1. Low TCO

Cloud service providers, multi-tenant colocation data centers and other businesses that emphasize operating efficiency should be sure that any UPS they put in their white space offers these essential features:

- **Modular architecture** – In the past, data centers in need of additional UPS capacity usually had to buy an entirely new UPS. Today's most advanced UPSs, however, feature multiple modular power units that share common electrical feeds, power supplies and cabling. Such products reduce the amount of time required to scale up in response to increased demand from days to hours.

Moreover, "growing vertically" by augmenting the capacity of existing UPSs rather than adding new ones also lowers hardware procurement costs, saves valuable floor space and provides needed flexibility for a growing cloud or colocation provider. Indeed, compared to the cost of purchasing, installing and maintaining a new UPS, adding 100 kW of capacity to an existing device instead can save companies up to \$35,000 in upfront expenses and another \$700 annually in footprint savings, based on yearly data center operating costs of \$150 per square foot.

Furthermore, some modular UPSs enhance reliability by providing internal redundancy. Should one power unit fail, the others around it can automatically take up the additional load. That enables cloud and colocation data centers to meet steep availability requirements without buying and maintaining expensive backup UPS hardware.

- **Industry-leading energy efficiency** – Power is a major operating expense for most cloud and colocation data centers, so such companies must pay close attention to energy efficiency when evaluating UPSs. The latest white space-ready models are 97 percent efficient in double-conversion mode and come with advanced operating modes capable of achieving 99.1 percent efficiency. Furthermore, the most sophisticated new white space UPSs maintain high efficiency levels across the entire operating load range, including light loads. That's an important consideration in data centers that typically load UPSs at less than 50 percent of capacity.



Figure 1. Modular UPSs save money and floor space by enabling data centers to "grow vertically" (reducing horizontal sprawl as shown above) and leverage internal redundancy, rather than install backup units.

2. Flexible deployment

Cloud computing providers and MTDC employ a variety of hardware deployment schemes, and reposition hardware frequently. They also aim to make optimal use of their floor space, for maximum profitability. When choosing a white space UPS, therefore, they should seek out products with these features:

- **Support for multiple configurations** – Certain cloud and colocation data centers position white space UPSs against the walls. Others place them within or at the end of server rows. Over time, moreover, new requirements can cause them to change their deployment preferences. Selecting UPSs capable of fitting comfortably in multiple configurations is therefore critical. Specifically, companies with distributed power architectures should utilize UPSs that provide:
 - **Multiple thermal management options**, including front-to-top airflow for use in wall and row configurations and front-to-back airflow for use in hot aisle/cold aisle configurations. Look as well for UPSs with capacity to add a narrow chimney that can be deployed within or at the end of a server row yet still leave room overhead for busway and cable trays.
 - **Line and match accessories** that fit standard rack depths and can be easily mirrored at either end of a row.
 - **A slim, line and match, detachable maintenance bypass** that eliminates the need for expensive conduit.
 - **Top and bottom cable entry**, for use with either slab or raised floors.





Figure 2. Multi-tenant data centers should use white space UPSs with both front-to-top and front-to-back airflow, to ensure they have the flexibility they need to support wall, aisle and row configurations.

3. Superior manageability

To honor the demanding commitments in their service level agreements, cloud and colocation vendors need white space hardware with a history of trouble free service and uncomplicated maintenance. Customers of MTDCs prefer maintenance windows that are relatively short. For providers with a distributed power architecture, that means selecting UPSs with these capabilities:

- **Front serviceability** – UPSs that offer full service access from the front of the unit increase uptime by ensuring that technicians can easily inspect or repair any internal component that requires their attention regardless of whether the unit is positioned against the wall, in an aisle or in a row.
- **Seamless integration with intelligent management software** – Companies with distributed power protection schemes can further enhance uptime and efficiency by enabling administrators to view and control everything in their white space—including UPS hardware—through a single centralized console. The most advanced new white space-ready UPSs support that goal by integrating tightly with leading management environments, such as Microsoft System Center and VMware vCenter Server. The management systems also provide access to detailed performance analytics that can help managers of fast-growing data centers optimize their power and cooling consumption patterns and quickly spot UPSs that are overloaded, likely to fail or in need of rebalancing.
- **24x7 remote monitoring and notification** – Cloud providers and MTDC operators need to know about trouble in their white space immediately, regardless of when it occurs. Companies should therefore seek out white space UPSs backed by round-the-clock monitoring services capable of providing in-progress notification of significant events via phone, email or pager.

Conclusion

Distributed power architectures that place UPSs in the white space can increase the scalability, availability and flexibility of MTDC and cloud services providers while lowering capital expenditures. To collect those benefits, however, companies must choose UPSs designed to reduce TCO, boost deployment flexibility and streamline management. Though previously unavailable, products with all of those qualities have finally reached market. MTDCs and cloud providers with centralized power schemes should therefore investigate switching to distributed power, while providers already using distributed power should consider implementing newer, more efficient white space UPSs.

About Eaton

Eaton's electrical business is a global leader with expertise in power distribution and circuit protection; backup power protection; control and automation; lighting and security; structural solutions and wiring devices; solutions for harsh and hazardous environments; and engineering services. Eaton® is positioned through its global solutions to answer today's most critical electrical power management challenges.

Eaton is a power management company with 2013 sales of \$22.0 billion. Eaton provides energy-efficient solutions that help our customers effectively manage electrical, hydraulic and mechanical power more efficiently, safely and sustainably. Eaton has approximately 101,000 employees and sells products to customers in more than 175 countries. For more information, visit www.eaton.com.

About the author

Jason Anderson is a Product Line Manager for Small to Mid-Size Data Center Solutions in Eaton's Critical Power Solutions Division in Raleigh, North Carolina. A 14-year veteran in the power systems industry, Anderson has experience in UPS systems new product design, technical sales and applications engineering, as well as working in power quality engineering for data centers and cable television. He may be reached at JasonSAnderson@Eaton.com.

For more information, please visit,
Eaton.com/pq/whitepaper

Eaton
1000 Eaton Boulevard
Cleveland, OH 44122
United States
Eaton.com

© 2014 Eaton
All Rights Reserved
Printed in USA
Publication No. WP153006EN
July 2014

EATON
Powering Business Worldwide

Eaton is a registered trademark of Eaton.

All other trademarks are property of their respective owners.